

(12) UK Patent Application (19) GB (11) 2 348 025 (13) A

(43) Date of A Publication 20.09.2000

(21) Application No 9906231.7

(22) Date of Filing 19.03.1999

(71) Applicant(s)
International Business Machines Corporation
(Incorporated in USA - New York)
Armonk, New York 10504, United States of America

(72) Inventor(s)
Colin Leonard Bird
John Brian Ibbotson

(74) Agent and/or Address for Service
IBM United Kingdom Limited
Intellectual Property Department, Mail Point 110,
Hursley Park, WINCHESTER, Hampshire, SO21 2JN,
United Kingdom

(51) INT CL⁷
G06F 17/30

(52) UK CL (Edition R)
G4A AUDB

(56) Documents Cited
WO 97/39406 A1
Dialog Access No 02142125 & Computergram
International, n335 1, pCGN02180013, 18 Feb 1998.

(58) Field of Search
UK CL (Edition Q) G4A AUDB
INT CL⁶ G06F 9/46 17/30 17/60
Online: COMPUTER, EPODOC, INSPEC, JAPIO, TDB,
WPI

(54) Abstract Title
Message broker providing a publish/subscribe service and method of processing messages in a publish/subscribe environment

(57) Provided is a message broker supporting a publish/subscribe service. The broker is adapted to identify, 120-150, characteristics of the content of messages received from publisher application programs and to use the identification of message content characteristics to determine, 160, which of a set of message handler modules to pass the message to. The selected message handler module processes the message and then one or more derived messages are sent, 200, to subscribers for which this message type or message content has been requested. One or more message handler modules are preferably multimedia processing modules and they are selected in response to identification of particular multimedia types within the message content. Processing by the handler module may be dependent, 180, on subscriber requirements. The broker may be distributed (Fig 2).

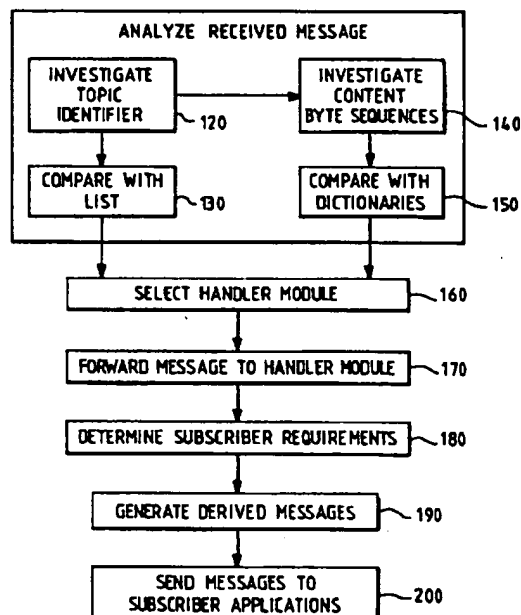


FIG. 3

GB 2 348 025 A



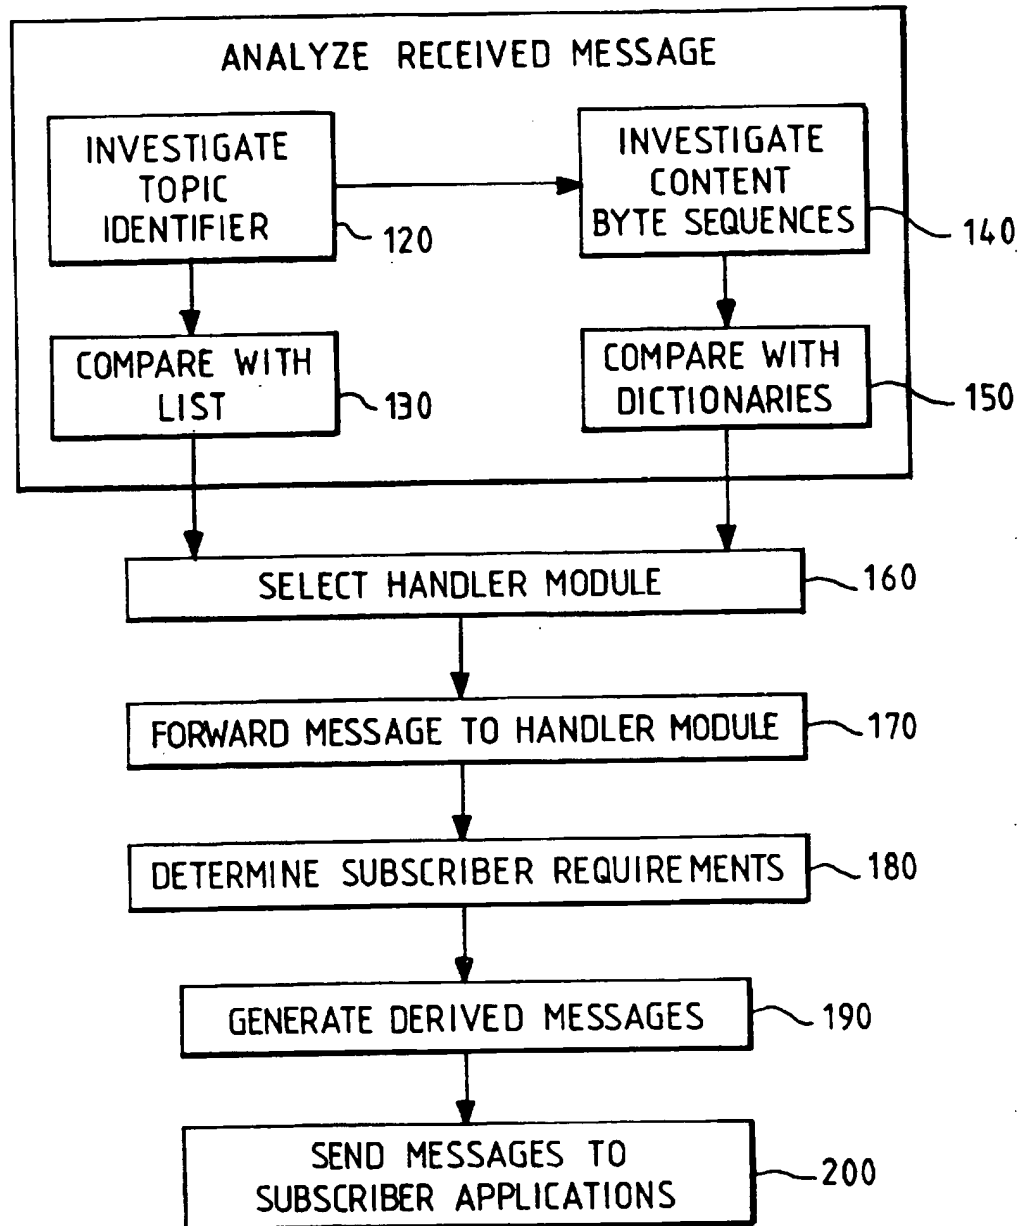


FIG. 1

FIG. 2

FIG. 3

MESSAGE BROKER PROVIDING A PUBLISH/SUBSCRIBE SERVICE AND
METHOD OF PROCESSING MESSAGES IN A PUBLISH/SUBSCRIBE ENVIRONMENT

Field of Invention

The present invention relates to a message broker providing a publish/subscribe service, and to a method of processing message content in a publish/subscribe environment.

Background

Publish/subscribe is a known information distribution mechanism in which information providers (Publishers) send information electronically across a network to a community of information users (Subscribers) who have chosen the topics of information they wish to receive. Typically, publishers publish a single message to a broker. The broker then decides from its registered subscriptions where to send copies or derivatives of the message.

Known publish/subscribe mechanisms rely on distribution lists, typically executing very simple logic to achieve required routing of highly structured, small transactional messages. Subscriber applications can be registered to receive all messages which include a specific subject classification or keyword, such as a company name or a more specific subject, within the message header or within a particular content field of the structured message. Message dictionaries hold information about the structure of messages, providing the message broker with a definition of the structure of all messages that will pass through it and so enabling unpacking of fields from the message content for inspection. The message broker queries the message header or a field within the content to retrieve the subject classification or keyword, and this is then compared with a list of subscriber applications' information requirements and any other stored rules for routing messages to determine which applications the message should be sent to.

Knowledge of the registered applications' required message formats also enables message format transformations (for example, an incoming message's content may be in COBOL whereas a subscriber may require this content to be converted to text within an e-mail memo).

These known publish/subscribe mechanisms can only operate with highly structured messages, since the simple logic involves comparison of the contents of only a specific message field with entries in a distribution list. Furthermore, known publish/subscribe mechanisms have typically only used simple logic for routing determination and format

transformations and have not extended the publish/subscribe environment to more complex processing of message content which is not text-based.

Summary of Invention

5

In a first aspect of the present invention, there is provided a message broker supporting a publish/subscribe service wherein the broker is adapted to identify characteristics of the content of messages received from publisher application programs and to use the identification of message content characteristics to determine which of a set of message handler modules to pass the message to. The selected message handler module processes the message and then one or more derived messages are sent to subscribers for which this message type or message content has been requested.

15

The function of identifying message content characteristics and logic for determining which of a set of message handler modules to pass the message to, and preferably also the processing functions of one or more message handler modules, may be implemented within a replaceable software component which interoperates with and controls the operation of a message broker. This separation of certain functions into a replaceable component (such as a Java Bean or ActiveX component) enables the behaviour of a message broker to be modified according to the particular software component which is used with the broker.

25

The message broker may comprise a distributed broker infrastructure, preferably with interoperating message broker software elements installed on each of a plurality of data processing systems. Messages are passed to different processing systems according to each system's capabilities and the identified message characteristics. For example, one processing system may be connected to a database to provide efficient SQL database access whereas a second system may have particular capabilities for image processing such as the ability to perform fast floating point operations. Thus, intelligent selection of message brokers of a distributed broker service can be achieved using the present invention's identification of message content characteristics.

35

In a second aspect, the invention provides a message broker supporting a publish/subscribe service wherein the broker is adapted to identify non-alphanumeric message content (i.e. content which is not text-based or numerical or a digital or other encoded representation of text or numbers) within messages received from publisher application programs and to use the identification of non-alphanumeric message content to determine which of a number of message handler modules to pass the message to. The selected message handler module processes the message

45

and then a derived message is sent by the broker to subscribers for which this message type or message content has been requested.

One or more of the selectable message handler modules are preferably multimedia processing modules, adapted to perform one or more specific processing operations on multimedia content within the message and then to pass one or more resulting modified or derived messages back to the broker. Recognition that a message contains non-alphanumeric content causes the broker to pass an incoming message to a processing module which is adapted to handle such content, such as a multimedia processing module. The processing module manipulates the content of the message and passes one or more derived messages back to the message broker, and it is the broker which then sends the derived message to the relevant subscriber application programs.

Thus, publish/subscribe capability has been extended by the present invention to support any type of message, including messages which include non-alphanumeric content such as multimedia content and to subscriber requirements which include a requirement for a message broker or associated message handler module to process the non-alphanumeric content. Prior to the present invention, identification and retrieval of any non-alphanumeric message content would have required a specific query to be constructed for each message scanned and so intelligent delivery of multimedia in a form required by subscribers could not be handled using known publish/subscribe mechanisms.

According to one embodiment of the invention, the message broker performs a determination of what particular operations are to be performed by the selected multimedia processing module according to specified requirements of subscribers and according to characteristics of the received message. Alternatively, a selected processing module may perform a standard set of operations regardless of subscriber requirements, with the broker only subsequently checking which subscribers require which type of output.

The non-alphanumeric content may be, for example, audio or image data or any binary data including executable program code. One example is where a broker is connected to receive messages containing waveform data representing the material broadcast by a radio station. The broker recognises that the message includes audio content and passes the message to a particular message handler module which performs the following functions: retrieves a watermark from the audio data, compares the watermark with a database of copyright owners and their material to identify the relevant owner and material, and generates a message to be published to the copyright owner informing them that a particular copyright work has been broadcast by a particular radio station at a

particular date and time. This will allow the copyright owner to seek royalty payments. A second example is where an incoming message includes raw image data captured by a satellite. The incoming message is processed to create sets of different resolution images, such that the appropriate resolution images are then available for delivery in response to different subscriber requests. Satellite images can also be processed in order that the published messages contain derived information such as crop classification maps.

Note that the above examples involve the processing modules significantly modifying or interpreting the received message content - the messages subsequently delivered to subscribers may include extracted, processed or derived data which is pertinent to the particular subscriber's requirements. There may also be examples where it is appropriate for messages sent to subscribers (referred to herein as "derived" messages) to include identical content to the messages published by publishers, as determined by the message content characteristics and/or subscriber requirements. The invention's identification of multimedia content and consequent selection of an appropriate handler module is equally applicable to these latter cases.

The message broker according to the preferred embodiment of the invention is not limited to a single, specific message structure or to scanning of the contents of a particular field within a structured message, but is adapted to scan the content of a received message and to identify non-alphanumeric elements within the message. For example, the broker may be adapted to look for a topic string with certain topics having an expectation of non-alphanumeric content, and then to look for recognisable byte sequences which are representative of the structure of particular image formats (for example GIF, TIFF and JPEG files have a recognisable initial byte sequence).

Brief Description of Drawings

Embodiments of the present invention will now be described in more detail, by way of example, with reference to the accompanying drawings in which:

Figure 1 shows a computer network including publisher and subscriber application programs exchanging messages via a message broker according to an embodiment of the present invention;

Figure 2 shows an example distributed computer network topology with message broker functions distributed across the network; and

Figure 3 shows the sequence of steps of a publish/subscribe method using a message broker according to an embodiment of the invention.

Detailed Description of Preferred Embodiments

5

Application programs running on different computers within a network are enabled to communicate and hence interoperate using messaging products such as IBM Corporation's MQSeries message queuing software products. Message queuing and commercially available message queuing products are described in "Messaging and Queuing Using the MQI", B.Blakeley, H.Harris & R.Lewis, McGraw-Hill, 1994, and in the following publications which are available from IBM Corporation: "An Introduction to Messaging and Queuing" (IBM Document number GC33-0805-00) and "MQSeries - Message Queue Interface Technical Reference" (IBM Document number SC33-0850-01). The network via which the computers communicate using message queuing may be the Internet, an intranet, or any computer network. IBM and MQSeries are trademarks of IBM Corporation.

IBM's MQSeries messaging software products provide transactional messaging support, synchronising messages within logical units of work in accordance with a messaging protocol which gives assured once and once-only message delivery even in the event of system or communications failures. MQSeries products provide assured delivery by not finally deleting a message from storage on a sender system until it is confirmed as safely stored by a receiver system, and by use of sophisticated recovery facilities. Prior to commitment of transfer of the message upon confirmation of successful storage, both the deletion of the message from storage at the sender system and insertion into storage at the receiver system are kept 'in doubt' and can be backed out atomically in the event of a failure. This message transmission protocol and the associated transactional concepts and recovery facilities are described in international patent application WO 95/10805 and US patent 5465328, which are incorporated herein by reference.

The message queuing inter-program communication support provided by the MQSeries products enables each application program to send messages to the input queue of any other target application program and each target application can asynchronously take these messages from its input queue for processing. This provides for assured delivery of messages between application programs which may be spread across a distributed heterogeneous computer network, but there can be great complexity in the map of possible interconnections between the application programs.

This complexity can be greatly simplified by including within the network architecture a communications hub to which all other systems connect, instead of having direct connections between all systems.

Message brokering capabilities can then be provided at the communications hub to provide intelligent message routing and integration of applications. Message brokering functions typically include the ability to route messages intelligently according to business rules and knowledge of different application programs' information requirements, using message 'topic' information contained in message headers, and the ability to transform message formats using knowledge of the message format requirements of target applications or systems. Such brokering capabilities are provided by IBM Corporation's MQSeries Integrator software product, providing services for messages which are exchanged between application programs using IBM's MQSeries messaging products.

If we adopt an hierarchical view of these various components with application programs at the top closest to the user, then a message broker component sits at the next layer down between the application programs and the message delivery software.

According to the present invention, a message broker software product or software component can be adapted to support the publish/subscribe mechanism for information delivery and to extend that mechanism to include more complex processing of message content.

Figure 1 shows an overview of message flows between publishing application programs 10 and subscribing application programs 20 via a message broker 30 according to the present invention. The message broker 30 according to the preferred embodiment of the invention is implemented as a software component which interoperates with a message delivery software product 90 running on a computer system 100 to which all publishing applications connect via a communications network. However, the brokering functions could be distributed across a plurality of message broker components connected via the network, as discussed in the next paragraph. Thus, the message broker (whether a single component or distributed) is connected to receive all published messages. Each application program is associated with a message delivery product 90' which handles the complexities of inter-program communication across the network, and each message broker is also associated with a message delivery product 90 on the local computer system.

As noted above, instead of a single broker at the communications hub, the hub may consist of a distributed collection of brokers 30 which communicate with each other. These brokers 30 may share subscription lists and provide a redundant, scalable distributed broker infrastructure. The message broker services provided by such a distributed broker infrastructure can be enhanced by using the present invention's identification of message content characteristics to select one of a plurality of processors within a distributed environment with

reference to its functional capabilities and the message characteristics such that processing of the message is optimised. For example, if one part of the message broker function is located on a system having a direct connection to a database, the efficiency of SQL queries may be optimised by passing all SQL queries to that system. If a second part of the message broker function is located on a system has specialised processing capabilities for images, then all messages containing images may be sent to that system.

For ease of understanding, the following description will refer to a single message broker located on a single data processing system to which all published messages are sent. When a message is received by the message broker, a message analysis component 60 within the broker firstly investigates the message header for a topic or subject identifier. A message publisher application may have identified a message topic or subject, such as by entering "STOCK/COMPUTERS/IBM" or "STOCK/AUTOS/GM" into a relevant field of the message header to indicate that the subject of the message is stock prices for IBM Corporation or General Motors respectively. A subscriber can specify subjects about which they wish to receive messages in a similar manner, such as by specifying "STOCK/COMPUTERS/*" for stock prices of all computer companies, or "*/IBM" for all messages which relate to IBM Corporation.

The message broker includes a rules engine 40 which compares any topic identifier retrieved from the message header with a list of subscribers' information requirements to identify which subscribers wish to receive this message. Additionally, the topic identifier is compared with a list of topic identifiers which are known to imply that there is some non-alphanumeric message content. In one example implementation, message publisher applications provide information to the broker (as part of a registration process) including a list of topic identifier codes which are to be used when the publisher is publishing messages including non-alphanumeric content.

Alternatively, the registration process may involve publishers specifying message topics which will only include text-based content, to indicate that no further analysis of the content of such messages is required to determine whether there is any non-textual content.

Secondly, unless the attempt to retrieve and interpret a message topic identifier enables a final determination of how to handle the message, the message broker's analysis component 60 scans the content of the message, looking for byte sequences which are representative of particular content types. The analysis component 60 compares byte sequences from the message content with dictionaries of byte sequences which are representative of specific content-type specific structures.

For example, GIF, TIFF and JPEG images each have characteristic byte structures. The content scanning step can identify audio content and images and resolve image formats. A message could also include program code, such as HTML (Hypertext Markup Language) or XML (Extensible Markup Language) codes which describe the format of data for display by an Internet Web browser. The message content scanning step could easily be implemented to recognise XML or HTML by reference to their respective dictionary of tags.

Having performed this investigation, any messages identified as having non-alphanumeric message content are passed 170 to an input buffer of an appropriate selected one of a set of multimedia processing modules 70, the module being selected 160 according to whether the non-alphanumeric content is image or audio information or some other data type.

Using the example of XML message content given earlier, examination of the XML content could indicate the existence of specific multimedia content and reveal the characteristics of the content that would be relevant to a determination of how it should be processed, and in some circumstances XML format could then also be used to provide one or more code modules that would be required for appropriately processing the identified content.

A collection of message handler modules 70, including the multimedia processing modules, are held as objects in an object library 80 and are instantiable in response to the broker's determination of a requirement for a specific processing module 70 to perform its processing operations on a received message. The message broker 30 responds to a selection 160 of a message handler module 70 by checking whether that module is already running (i.e. has been instantiated on receipt of a previous message) and, if so, queuing the message in an input buffer for processing by that handler module. If the required handler module is not running, then it is instantiated and the message is passed to the input buffer of the handler module for processing.

In one simple embodiment of the invention, all messages which include image content are passed to a single image processor module. Similarly, all messages including audio data are passed to a single audio processing module and messages including other identifiable data types are passed to a respective processing module for that type. This may entail passing a single message through a sequence of processing steps each performed by a different message handler module if the message includes a number of content types. Messages containing executable code may be passed to a processing module that performs authentication (for example, digital signature verification or virus checking) prior to

permitting publication, to protect subscribers from undesirable program code. This provision of a single message handler module for each type of message content means that an appropriate message handler module can be selected by a relatively simple analysis of the incoming message.

5

In one embodiment of the invention, a selected message handler module then performs a generic set of processing operations - i.e. a set of identical operations are performed on all messages which are passed to that handler module. However, this is not an optimal implementation for many message handler modules.

10

In a preferred embodiment of the invention, the processing which is performed by the selected message handler module is determined by the specified requirements of individual subscriber applications and the particular characteristics of the received message. Thus, following selection of a message handler module, the message broker determines which subscriber application programs have specified a requirement to receive a published message in response to the particular topic or type of the received message and then each subscriber's requirements are used to generate a task description which is input to the selected message handler module together with the message.

15

20

The selected message handler module then performs 190 operations on the message content in accordance with the associated task description, and passes the resulting derived message or messages back to the message broker for distribution to the relevant subscribers.

25

In yet a further embodiment of the invention, the actual selection of message handler modules is itself dependent on the specified requirements of subscriber application programs, which means that the reference to subscriber requirements is made by the broker as part of the selection step instead of afterwards.

30

The capabilities of the invention will now be described with reference to a number of examples of message processing and brokering according to the invention.

35

Consider a computer system which is set up to publish messages including images of the surface of the earth captured by an orbiting satellite. The organisation obtaining the information is able to share the costs of retrieving and processing their data by offering a subscription service to other organisations interested in receiving particular data contained within the captured images. A publisher application program at this first system periodically generates messages containing the raw captured images and sends these messages to the message broker. The message broker analyses 120-150 the messages on

40

45

receipt and identifies the presence of images within the messages (by the combination of comparing any recognised message topic identifier with known topics and scanning message content, as described above).

5 The message broker determines 160, in response to recognition of an image, that the message should be passed to an image processing message handler module. The broker then uses the topic identifier to check 180 records of subscriber requirements - determining which subscribers have registered a requirement to receive messages related to this topic and
10 determining what specific data requirements have been specified by those subscribers. As an example, a first subscriber may have registered with the message broker with a requirement for the subscriber to be set a message whenever the message broker receives messages including a topic identifier naming the specific publisher organisation. The subscriber may
15 also have specified that all received messages from that publisher organisation should have their image content processed to generate a derived image showing classifications of crops identifiable from the captured raw image. Additionally, the subscriber may have specified an output type requiring processing to significantly transform the message
20 into the required type - for example, if a subscriber application program is connected to an audio output device it may require all messages to be converted to a waveform output file whereas another subscriber application program may be an e-mail program requiring messages to be sent as e-mail memos.

25 A task description is generated which describes the operations to be performed to process the received message's image content in the required way, and the task description and received message are then passed 170 to the selected image processing message handler module. This
30 message handler module performs 190 the described sequence of processing operations and then passes back to the message broker a new message which includes the new processed image data.

35 The message broker then forwards 200 the new message to the subscriber or subscribers who specified this information requirement. This routing of the final derived message or messages to subscribers in accordance with subscriber's information requirements (under the control of the rules engine 40) is very similar to the routing performed by prior art message brokers. Thus, the present invention enables content-type
40 specific processing of message content and the ability to perform appropriate processing of a message without reliance on a single specific message structure to be incorporated within the context of an otherwise conventional publish/subscribe system which uses business logic for identification of relevant subscribers.

45

In the preferred embodiment of the invention, the message content type is used to select a message handler module without reliance on a topic identifier in the message header. The analysis component 60 of the message broker scans message content (and this can be done either in all cases or only if querying of the message header is not sufficient to determine which message handler module to select). When an image is identified within the message content by this scanning step, the image can undergo a watermark extraction process followed by an analysis of the watermark to determine whether this includes information sufficient to identify what message handler module should be selected or what specific processing operations should be performed on the image. This may enable performance of the handler selection stage, and possibly also the message content processing stage, of implementation of the invention prior to investigating specific subscriber requirements.

Similar watermark extraction and analysis to determine required processing operations can also be implemented for audio data within a message, as discussed previously.

The message broker according to the preferred embodiment of the invention is also able to employ logic which relies on an analysis of both text-based information in a message header and text-based data within the content of a message to determine whether a particular subscriber should receive a message in response to receipt by the broker of an incoming message. That is, in addition to the multimedia processing aspect of the invention, the text processing performed by a broker according to the invention may include processing subscriber information requirements such as a requirement to be notified of the stock price of a company (COMPANYNAME) only when the stock price exceeds a threshold price. The subscriber may specify "STOCK/COMPUTERS/COMPANYNAME" as a topic identifier to be identified from message headers and "PRICE: >\$150" as a filter which requires scanning of a share price field within the content of a received message to determine whether a message should be sent to that subscriber. The content analysis can also be implemented by queries which use the Structured Query Language notation or similar to test whether certain conditions are true (for example, "the value of COMPANYNAME stock is > \$150").

In a further example of operation of a message broker implementing the invention, content queries may involve multimedia content analysis instead of the simple business logic which is known in prior art systems. For example, subscribers may register an interest in receiving any images which include a red circle towards the top right of the image as a way of investigating possible trade mark infringement, or register an interest in receiving any image which is sufficiently similar to an example image or set of images which are representative of certain classes of image.

Technology is available (for example, IBM's Query By Image Content (QBIC) technology described in US patent 5579471) which allows image matching in response to such example images or user descriptions of features of an image that they want to match. The incoming image would be run as a query against a QBIC catalogue comprising the class representatives, the results being a set of distance measures. Fuzzy matching algorithms establish how close an image within a message is to the examples images or requested features. The subscription registration mechanism in a traditional publish/subscribe broker can thus be enhanced using multimedia feature definitions and multimedia feature comparisons to allow only those messages that meet the subscription criteria (within a given tolerance defined by specific thresholds) to be delivered to the subscribing user.

The QBIC example shows that the identification of message content characteristics can comprise separate steps of recognising the existence of a certain message content type (typically using precise matching of a given byte sequence) and then identifying image content characteristics beyond those which are inherent in the recognised byte sequence (such as whether the content matches defined parameters or example images within certain tolerances).

Following these steps, a decision is made as to what processing to perform and how to do it. There is considerable flexibility within the scope of the invention regarding how to implement this step. The decision can employ a stored and relatively inflexible method or a dynamic approach that relies on input from both publisher and subscriber. The subscriber may specify when registering as a subscriber some particular processing operations that it requires to be performed on messages to generate derived messages. Examples of this were described earlier.

As noted previously, the function of identifying message content characteristics and the logic for determining which of a set of message handler modules to pass the message to, and preferably also the processing functions of one or more message handler modules, may be implemented within a replaceable software component which interoperates with and controls the operation of a message broker. This separation of certain functions into a replaceable component (such as a Java Bean or ActiveX component) enables the behaviour of a message broker to be modified according to the particular software component which is used with the broker. It also enables new capabilities to be added over time, such as by loading new message feature dictionaries into an object library when they become available and thereby enabling the message broker to recognise new message content types.

CLAIMS

1. A message broker for supporting a publish/subscribe service by receiving messages from publisher application programs and then forwarding messages to subscriber application programs in accordance with subscriber requirements, wherein the message broker component includes:

a message analysis component for identifying non-alphanumeric message content within messages received from publisher application programs;

a selector for using the identification of non-alphanumeric message content to determine which of a set of message handler components to pass the message to;

a set of selectable message handler components for processing received messages to generate one or more derived messages;

means for accessing a repository storing requirements of subscriber application programs; and

means, responsive to requirements of subscriber application programs stored in said repository, for passing derived messages to subscriber application programs which require the derived messages.

2. A message broker according to claim 1, which is adapted to access the repository of subscriber requirements to determine what processing operations are to be performed on received messages by the message handler components.

3. A message broker according to claim 1 or claim 2, wherein the selector is adapted to select one of the set of message handler modules in response to an identification by the message analysis component of the type of non-alphanumeric message content within a received message.

4. A message broker according to claim 3, wherein the means for identifying non-alphanumeric message content includes means for analysing byte sequences within the message content to identify structures representative of particular types of multimedia message content.

5. A message broker according to any one of the preceding claims, wherein the message analysis component is adapted to implement a watermark retrieval operation for received messages and to analyse retrieved watermarks to identify information contained therein which facilitates a determination of which of the set of message handler components to pass the message to.

6. A message broker according to any one of claims 1 to 4, wherein one or more of said message handler modules is adapted to implement a watermark retrieval operation for messages passed to that message handler module and to analyse retrieved watermarks to identify information contained therein which facilitates a determination of what processing operations are to be performed on the message by the message handler module.

7. A computer program product comprising computer readable program code stored on a computer readable storage medium, the program code comprising a message broker according to any one of claims 1 to 6.

8. A process controller component for a message broker according to claim 1, wherein the process controller component includes the message analysis component and the message handler component selector.

9. A process controller component according to claim 8, including at least one message handler component of said set of selectable message handler components for processing received messages to generate derived messages.

10. A message broker software component supporting a publish/subscribe service, for controlling the operation of a computer to execute processes

to identify non-alphanumeric content of messages received from publisher application programs,

to use the identification of non-alphanumeric message content to determine which of a set of message handler modules to pass the message to,

to process the message using the selected message handler module to generate one or more derived messages, and

to send a derived message to subscribers which require the derived messages, as determined by stored subscriber requirements.

11. A distributed message broker infrastructure comprising a plurality of message broker software components according to claim 10, each for running on one of a plurality of interconnected computer systems, each message broker software component having access to information regarding the capabilities of each of the plurality of interconnected computer systems and each message broker software component being adapted to select, in accordance with the identification of non-alphanumeric content of received messages and in accordance with the capabilities of the

plurality of interconnected computer systems, one of said computer systems for processing a message.

5 12. A method for sending messages to subscriber application programs in response to receipt of messages from publisher application programs, the method including the following steps performed under the control of a message broker:

10 identifying non-alphanumeric content of messages received from publisher application programs;

determining which of a set of message handler modules to pass the message to, using the identification of non-alphanumeric message content;

15 processing the message using the selected message handler module to generate one or more derived messages; and

20 sending derived messages to subscribers which require the derived messages, as determined by stored subscriber requirements.

25 13. A method, implemented by a message broker within a communications network, of processing messages received by the message broker from publisher application programs and publishing messages to subscriber application programs in response to said received messages, the method including the steps of:

analysing received messages to identify non-alphanumeric message content;

30 selecting, using the identification of non-alphanumeric message content, one of a set of message handler components to pass each received message to;

35 passing each received message to the selected message handler module;

40 performing the steps of retrieving information requirements of subscriber application programs from a repository storing said information requirements and, either with or without reference to said retrieved information requirements, processing the received messages using the selected message handler modules; and

45 responsive to the retrieved information requirements of subscriber application programs, passing processed messages to subscriber application programs which have a requirement to receive the processed messages.

14. A method according to claim 13, wherein the communications network is the Internet.

5 15. A computer system including a message broker component for supporting a publish/subscribe service on behalf of application programs connected thereto via a communications network, wherein the message broker component includes:

10 a message analysis component for identifying non-alphanumeric message content within messages received from publisher application programs;

15 a selector for using the identification of non-alphanumeric message content to determine which of a set of message handler modules to pass the message to;

a set of selectable message handler components for processing received messages to generate one or more derived messages;

20 a repository for storing requirements of subscriber application programs; and

25 means, responsive to requirements of subscriber application programs stored in said repository, for passing derived messages to subscriber application programs which require the derived messages.



Application No: GB 9906231.7
Claims searched: 1 to 15

Examiner: Geoff Western
Date of search: 21 October 1999

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.Q): G4A AUDB

Int Cl (Ed.6): G06F 9/46 17/30 17/60

Other: Online : COMPUTER, EPODOC, INSPEC, JAPIO, TDB, WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	WO 97/39406 A1 (NEW ERA SOFTWARE)	-
A	Dialog Access No 02142125 & Computergram International, n3351, pCGN02180013, 18 Feb 1998, "Tibco adds message broker, trumpets enterprise cred".	-

X Document indicating lack of novelty or inventive step
Y Document indicating lack of inventive step if combined with one or more other documents of same category.
& Member of the same patent family

A Document indicating technological background and/or state of the art.
P Document published on or after the declared priority date but before the filing date of this invention.
E Patent document published on or after, but with priority date earlier than, the filing date of this application.